

MAXIMUM LIKELIHOOD ESTIMATION IN LOG-LINEAR MODELS¹

BY STEPHEN E. FIENBERG AND ALESSANDRO RINALDO

Carnegie Mellon University

We study maximum likelihood estimation in log-linear models under conditional Poisson sampling schemes. We derive necessary and sufficient conditions for existence of the maximum likelihood estimator (MLE) of the model parameters and investigate estimability of the natural and mean-value parameters under a nonexistent MLE. Our conditions focus on the role of sampling zeros in the observed table. We situate our results within the framework of extended exponential families, and we exploit the geometric properties of log-linear models. We propose algorithms for extended maximum likelihood estimation that improve and correct the existing algorithms for log-linear model analysis.

1. Introduction. Log-linear models are arguably the most popular and important statistical models for the analysis of categorical data; see, for example, Bishop, Fienberg and Holland (1975), Christensen (1997). These powerful models, which include as special cases graphical models [see, e.g., Lauritzen (1996)] as well as many logit models [see, e.g., Agresti (2002), Bishop, Fienberg and Holland (1975)], have applications in many scientific areas, ranging from social and biological sciences, to privacy and disclosure limitation problems, medicine, data-mining, language processing and genetics. Their popularity has greatly increased in the last decades due to growing demands for analyzing databases taking the form of large and sparse contingency tables, where most of the cell entries are very small or zero counts. Despite the widespread usage of these models, the applicability and statistical properties of log-linear models under sparse settings are still very poorly understood. As a result, even though high-dimensional sparse contingency tables constitute a type of data that is common in practice (e.g., in sam-

Received September 2011; revised February 2012.

¹Supported in part by NSF Grant DMS-06-31589.

AMS 2000 subject classifications. Primary 62H17; secondary 62F99.

Key words and phrases. Extended exponential families, extended maximum likelihood estimators, Newton–Raphson algorithm, log-linear models, sampling zeros.

This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in *The Annals of Statistics*, 2012, Vol. 40, No. 2, 996–1023. This reprint differs from the original in pagination and typographic detail.

ple survey applications), their analysis remains exceptionally difficult; see Erosheva, Fienberg and Joutard (2007) for such an example.

In this article we are concerned with statistical inference in log-linear models of arbitrary dimension, and, in particular, with conditions for the existence of the maximum likelihood estimator, or MLE, of the model parameters. In log-linear model analysis, virtually all methodologies for assessment of fit, model selection and interpretation are applicable and have theoretical validity only provided that the MLE exists. Though this may appear to be only a computational issue, in fact, when MLE is not defined, the applicability of statistical procedures routinely used by practitioners may no longer have a theoretical justification and, at the very least, require alteration. The statistical implications of a nonexistent MLE, some of which are detailed below, are numerous and severe.

- Existence of the MLE is required to justify the use of large sample χ^2 approximations to numerous measures of goodness-of-fit commonly utilized for model assessment and model selection; see, for example, Bishop, Fienberg and Holland (1975), Agresti (2002), Read and Cressie (1988). When the MLE does not exist, the standard regularity conditions used to derive such approximations no longer hold. As we show below, under a nonexistent MLE, the model is not identifiable, the asymptotic standard errors are not well defined and the number of degrees of freedom becomes meaningless. Though existence of the MLE is by no means enough to warrant the use of χ^2 approximations, nonexistence will surely make them inadequate.
- Existence of the MLE is also needed to derive a limiting distribution for the double-asymptotic approximations of the likelihood ratio and Pearson's χ^2 statistic for tables in which both the sample size and the number of cells are allowed to grow unbounded, a setting studied, among others, by Morris (1975), Haberman (1977) and Koehler (1986); see also Read and Cressie (1988).
- The issue of nonexistence is also important for Bayesian analysis of log-linear models; see, for example, King and Brooks (2001), Massam (2009), Dobra and Massam (2010) and references therein. Indeed, we will demonstrate that nonexistence of the MLE is due to the data not being fully informative about the model parameters, and results in nonestimability of those parameters. Since the nonexistence of MLEs is due to insufficient data, it cannot be remediated. In particular, the use of Bayesian methods in cases in which the MLE is nonexistent is equivalent to replacing the information content lacking in the data with the information contained in the prior. Since for some parameters no learning from the data takes place, the posterior distribution must be interpreted accordingly. Furthermore, when one uses improper priors for the log-linear parameters, the posterior may be also be improper when the MLE does not exist; see Forster (2004).

It has long been known [see, in particular, Birch (1963), Haberman (1974), Bishop, Fienberg and Holland (1975)] that the nonexistence of the MLE is caused by sampling zeros. When certain patterns of zero counts occur in the observed table, the log-likelihood function cannot be maximized by any vector of finite norm. While for hierarchical log-linear models, patterns of sampling zeros leading to null margins are well known to cause nonexistence of the MLE, very little has been known or observed about general patterns of sampling zeros associated with nonexistent MLEs. The very few known examples described in Haberman (1974), Fienberg and Rinaldo (2007) and Dobra et al. (2009) suggest that nonexistence of the MLE may occur in small tables, but is very likely to arise when the table is large and sparse.

Haberman (1974) first obtained necessary and sufficient conditions for the existence of the MLE for log-linear models. Eriksson et al. (2006) gave a direct geometric interpretation of Haberman's conditions and proposed a polynomial time algorithm for checking for the existence of the MLE. Aickin (1979) and Verbeek (1992) refined Haberman's conditions by recasting the problem within the frameworks of exponential families and of generalized linear models, respectively. In fact, the issue of nonexistence of the MLE is best dealt with using the general theory of exponential families and, in particular, of extended exponential families, originally put forward by Barndorff-Nielsen (1978) and then Brown (1986). See also the important work by Čencov (1982). In a recent series of papers, Csizsár and Matúš (2001, 2003, 2005, 2008) broadened significantly the notions of extended exponential families and extended maximum likelihood estimation to include very general settings under minimal assumptions. See, in particular, Remark 5.9 in Csizsár and Matúš (2008), which briefly point to the connections with the theory of log-linear models. Rinaldo, Fienberg and Zhou (2009) and Geyer (2009) contain more specialized results directly relevant to the log-linear settings. Adopting a different approach, Lauritzen (1996) *defined* the parameter space for log-linear models as the point-wise limit closure of the log-linear model parameters, which he calls the extended log-linear model, and effectively treats the MLE and extended MLE as one entity. While this is theoretically convenient, the issue of nonestimability of the model parameters is not resolved, and the computation of the extended MLE is just as problematic. Finally, Nardi and Rinaldo (2012) provided asymptotic conditions under which, for a hierarchical log-linear model, a penalized maximum likelihood estimator based on the group-lasso penalty will return the correct model, with high probability.

Despite the breadth of the cited literature, two key issues concerning maximum likelihood estimation in log-linear models remain. First, the properties of extended exponential families have not yet been specialized to the case of log-linear models. In particular, direct application of this theory does not yield, in general, usable conditions for the existence of the MLE, and the

identification of the nonestimable log-linear parameters or of the patterns of zeros leading to a nonexistent MLE are still open problems. Secondly, existing theoretical results have not been incorporated yet in any numerical algorithm for checking for existence of the MLE and for identifying nonestimable parameters. Consequently, virtually all statistical software currently available to practitioners is flawed, to the point that nonexistence of the MLE can be detected only by monitoring whether the algorithm used to optimize the log-likelihood function fails to converge, or converges slowly or becomes unstable; see, for example, Fienberg and Rinaldo (2007). Consequently, results and decisions stemming from the statistical analysis of contingency tables containing substantial numbers of zero counts can be seriously compromised.

In this article we attempt to rectify these problems. Our contributions are two-fold:

- From a theoretical standpoint, we derive necessary and sufficient conditions for existence of the MLE that are broadly applicable to a variety of sampling schemes and amenable to computations. Ultimately, these conditions amount to checking whether the observed sufficient statistics lie on the boundary a polyhedral cone, called the *marginal cone*; see Eriksson et al. (2006). When the MLE does not exist, we specialize the theory of extended exponential families to characterize the estimability of the natural and mean-value parameters of the log-linear models. To this end, we focus on discrete exponential families with polyhedral convex support [see Rinaldo, Fienberg and Zhou (2009), Geyer (2009)], and rely significantly on tools from polyhedral geometry.
- From a practical viewpoint, we develop algorithms for extended maximum likelihood estimation that are applicable to large tables. Our procedures will allow one to (i) detect nonexistence of the MLE and (ii) identify and estimate all the parameters that are in fact estimable. Overall, our algorithms correct and improve over many existing software for log-linear model analysis. Due to space constraints, a detailed description of these algorithms is contained in the supplementary material [Fienberg and Rinaldo (2012)].

Notation. We let \mathcal{I} be a finite set of indices or cells, representing the support of a discrete distribution, such as the joint distribution of a set of categorical variables. We set $I = |\mathcal{I}|$, where $|B|$ is the cardinality the set B . We denote by $\mathbb{R}^{\mathcal{I}}$ be the vector space of real-valued functions on \mathcal{I} , and $\mathbb{R}_{\geq 0}^{\mathcal{I}}$ and $\mathbb{N}^{\mathcal{I}}$ its subset of nonnegative functions and nonnegative integer-valued functions, respectively. For vectors \mathbf{x} and \mathbf{y} , $(\mathbf{x}, \mathbf{y}) = \mathbf{x}^{\top} \mathbf{y}$ represents their inner product and $\|\mathbf{x}\| = \sqrt{\mathbf{x}^{\top} \mathbf{x}}$ the corresponding Euclidean norm. If $\mathbf{x} \in \mathbb{R}^{\mathcal{I}}$, we denote by $\mathbf{x}(i)$ the value corresponding to the i th coordinate of \mathbf{x} and

by $\text{supp}(\mathbf{x}) = \{i : \mathbf{x}(i) \neq 0\}$ the set of coordinates of \mathbf{x} with nonzero values. We take functions and relations on vectors component-wise, for example, for $\mathbf{x} \in \mathbb{R}^{\mathcal{I}}$, $\exp(\mathbf{x}) = \{e^{\mathbf{x}(i)} : i \in \mathcal{I}\}$.

For a nonempty subset $\mathcal{F} \subseteq \mathcal{I}$, we let $\pi_{\mathcal{F}} : \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}^{\mathcal{F}}$ the coordinate projection map given by $\{\mathbf{x}(i) : i \in \mathcal{I}\} \mapsto \{\mathbf{x}(i) : i \in \mathcal{F}\}$ and, for any $S \subset \mathbb{R}^{\mathcal{I}}$, we set $\pi_{\mathcal{F}}(S) = \{\pi_{\mathcal{F}}(\mathbf{x}), \mathbf{x} \in S\}$. If \mathcal{M} is a linear subspace, we denote by \mathcal{M}^{\perp} its orthogonal complement and by $\Pi_{\mathcal{M}}$ the orthogonal projector into \mathcal{M} . If \mathcal{N} is another linear subspace contained in \mathcal{M} , we write $\mathcal{M} \ominus \mathcal{N}$ for the subspace $\mathcal{M} \cap \mathcal{N}^{\perp}$.

For a matrix A , $\mathcal{R}(A)$ denotes its column range and $\text{kernel}(A)$ its null space. If the rows of A are indexed by \mathcal{I} , and \mathcal{F} is a nonempty subset of \mathcal{I} , $A_{\mathcal{F}}$ is the submatrix of A comprised of the rows with indexes in \mathcal{F} . We write $\text{cone}(A)$ for the polyhedral cone spanned by columns of A and $\text{conv}(A)$ for the polytope consisting of the convex combinations of its columns. Similarly, for a set S , $\text{conv}(S)$ is the convex hull of all its points. For a polyhedron P , we write its relative interior as $\text{ri}(P)$.

2. Log-linear models, sampling schemes and exponential families. Log-linear model analysis is concerned with the study of discrete probability distributions over a finite set \mathcal{I} , whose elements will be referred to as *cells*. These distributions are assumed to form an exponential family of probabilities $\{P_{\boldsymbol{\eta}}, \boldsymbol{\eta} \in \mathbb{R}^d\}$ with densities with respect to the counting measure on \mathcal{I} of the form

$$(1) \quad p_{\boldsymbol{\eta}}(i) = P_{\boldsymbol{\eta}}(\{i\}) = \exp\{(\boldsymbol{\eta}, \mathbf{a}_i) - \phi(\boldsymbol{\eta})\}, \quad \boldsymbol{\eta} \in \mathbb{R}^d,$$

where each \mathbf{a}_i is a nonzero vector in \mathbb{R}^d , and $\phi(\boldsymbol{\eta}) = \log(\sum_i \exp\{(\boldsymbol{\eta}, \mathbf{a}_i)\})$ is the log-partition function. The $I \times d$ matrix A , whose i th row is the vector \mathbf{a}_i^{\top} , is called the *design matrix*.²

Suppose we observe a sample of N independent and identically distributed realizations from an unknown distribution satisfying (1), where the data take the form of an unordered sequence of random cells (L_1, \dots, L_N) , with $L_j \in \mathcal{I}$ for each j , and where N too can be random. The observed cells are then cross-classified into a random integer vector $\mathbf{n} \in \mathbb{N}^{\mathcal{I}}$, called a *contingency table*, with $\mathbf{n}(i) = |\{j : L_j = i\}|$, for all $i \in \mathcal{I}$.

Traditionally, log-linear model analysis is not directly concerned with the natural parameters $\boldsymbol{\eta}$ in (1), but rather with the unknown expected value $\mathbf{m} := \mathbb{E}[\mathbf{n}]$ of the resulting contingency table, under the provision that $\mathbf{m}(i) > 0$ for each i . In detail, letting $\mathcal{M} \subset \mathbb{R}^{\mathcal{I}}$ be the linear subspace spanned by the rows of the design matrix A , the ensuing log-linear model is predicated

²It is easy to see that design matrices are not uniquely determined: if A_1 and A_2 are two matrices of dimensions $I \times d_1$ and $I \times d_2$, respectively, and with identical row spans, then they parametrize the same statistical model.

on the condition that $\boldsymbol{\mu} := \log(\mathbf{m}) \in \mathcal{M}$. In particular, log-linear models are typically defined as statistical models for the distribution of the random table \mathbf{n} indexed by the points in the linear subspace \mathcal{M} .

The distribution of the table \mathbf{n} depends on the sampling scheme used during the data collection process. In this article, we study sampling schemes based on linear restrictions on \mathbf{n} , known as conditional Poisson sampling schemes, introduced in Haberman (1974), Chapter 1. Specifically, let \mathcal{N} be a given m -dimensional linear subspace of \mathcal{M} , which we will refer to as the *sampling subspace*, and \mathbf{c} a known vector in $\mathbb{R}^{\mathcal{I}}$. The corresponding conditional sampling Poisson scheme prescribes that the distribution of \mathbf{n} is given by the conditional distribution of I independent Poisson random variable $\{\mathbf{n}(i), i \in \mathcal{I}\}$ with mean parameters $\{\mathbf{m}(i) = \exp(\boldsymbol{\mu}(i)), i \in \mathcal{I}\}$, where $\boldsymbol{\mu} \in \mathcal{M}$, given that $\Pi_{\mathcal{N}} \mathbf{n} = \mathbf{c}$. This type of data sampling includes the most commonly used sampling schemes, described below.

- *Poisson sampling scheme.* The sampling subspace is $\mathcal{N} = \{\mathbf{0}\}$. Thus, there are no restrictions on \mathbf{n} , which is a random vector comprised of independent Poisson random variables with mean \mathbf{m} . The log-likelihood function is given by

$$(2) \quad \ell^P(\boldsymbol{\mu}) = (\mathbf{n}, \boldsymbol{\mu}) - \mathbf{1}^\top \exp(\boldsymbol{\mu}) - \sum_i \log \mathbf{n}(i)!, \quad \boldsymbol{\mu} \in \mathcal{M}.$$

- *Product multinomial and multinomial sampling schemes.* Let $\mathcal{B}_1, \dots, \mathcal{B}_m$ be a partition of \mathcal{I} . Under the product multinomial sampling, the conditional distribution of the cell counts \mathbf{n} is the product of m independent multinomials of sizes N_j , $j = 1, \dots, m$, each supported on the corresponding class \mathcal{B}_j . Formally, let χ_j be the indicator function of \mathcal{B}_j , where $\chi_j(i)$ is 1 if $i \in \mathcal{B}_j$ and 0 otherwise, and define \mathcal{N} to be the r -dimensional subspace spanned by the orthogonal vectors (χ_1, \dots, χ_r) . The product multinomial sampling constraints are of the form $(\mathbf{n}, \chi_j) = N_j$, for known integer constants N_j . The log-likelihood function is [see Haberman (1974), equation 1.51]

$$(3) \quad \tilde{\ell}^M(\boldsymbol{\mu}) = \sum_{j=1}^r \left(\sum_{i \in \mathcal{B}_j} \mathbf{n}(i) \log \frac{\mathbf{m}(i)}{(\mathbf{m}, \chi_j)} + \log N_j! - \sum_{i \in \mathcal{B}_j} \log \mathbf{n}(i)! \right),$$

$\boldsymbol{\mu} \in \mathcal{M},$

where $\mathbf{m} = \exp(\boldsymbol{\mu})$. Because of the sampling constraints, $\tilde{\ell}^M$ is well defined only on the subset of \mathcal{M} ,

$$(4) \quad \widetilde{\mathcal{M}} := \{\boldsymbol{\mu} \in \mathcal{M} : (\chi_j, \exp(\boldsymbol{\mu})) = N_j, j = 1, \dots, r\},$$

which is neither a vector space nor a convex set. We give a more convenient parametrization below in Lemma 2. The multinomial scheme is

a special case of product multinomial schemes, corresponding to the trivial one-class partition of \mathcal{I} with indicator function $\mathbf{1}$. In this case, \mathbf{n} has a multinomial distribution with size $N = (\mathbf{1}, \mathbf{n}) = (\mathbf{1}, \mathbf{m})$ and cell probabilities \mathbf{m}/N .

- *Poisson–multinomial sampling schemes.* This sampling scheme is a combination of the previous two schemes. For a given partition $\mathcal{B}_1, \dots, \mathcal{B}_m$ of \mathcal{I} , the sampling constraints are of the form $(\mathbf{n}, \boldsymbol{\chi}_j) = N_j$ for $j = 1, \dots, m-1$, with the counts for the cells in the set \mathcal{B}_m left unconstrained; see [Lang \(2004, 2005\)](#).

As is customary, we assume throughout that the sampling subspace \mathcal{N} is strictly contained in \mathcal{M} . The case $\mathcal{N} = \mathcal{M}$ is practically uninteresting, as the resulting sampling constraints would fix the value of the sufficient statistics so that the conditional distribution of \mathbf{n} will not depend on the model parameters. We treat the case $\mathcal{N} \not\subset \mathcal{M}$ in the supplementary material [Fienberg and Rinaldo (2012)].

We now derive the equivalent exponential family representation for log-linear models under conditional Poisson schemes. To this end, we will express the sampling constraints in a different, but equivalent, form. Let $(\mathbf{v}_1, \dots, \mathbf{v}_m)$ be any set of m vectors spanning \mathcal{N} and such that $(\mathbf{v}_j, \mathbf{c}) = 1$ for all j . Then, the sampling constraints take the form

$$\mathbf{V}^\top \mathbf{n} = \mathbf{1},$$

where \mathbf{V} is the $I \times m$ matrix whose j th column is \mathbf{v}_j . Accordingly, we denote with

$$S(\mathbf{V}) := \{\mathbf{x} \in \mathbb{N}^{\mathcal{I}} : \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}$$

the set of all possible tables compatible with the sampling constraints specified by \mathbf{V} . Let ν be the finite measure on $\mathbb{N}^{\mathcal{I}}$ given by³

$$\nu(\mathbf{x}) := \prod_{i \in \mathcal{I}} \frac{1}{\mathbf{x}(i)!}, \quad \mathbf{x} \in \mathbb{N}^{\mathcal{I}}.$$

For a conditional Poisson scheme defined by \mathbf{V} , let $\nu_{\mathbf{V}}$ be the restriction of ν on $S(\mathbf{V})$, that is, $\nu_{\mathbf{V}}(\mathbf{x}) := 1_{\mathbf{x} \in S(\mathbf{V})} \nu(\mathbf{x})$, with $\mathbf{x} \in \mathbb{N}^{\mathcal{I}}$.

It is easy to see that the conditional distribution of the table \mathbf{n} , given the sampling constraints determined by \mathbf{V} , is the exponential family of distributions with base measure $\nu_{\mathbf{V}}$, sufficient statistics $\mathbf{A}^\top \mathbf{x}$, natural parameter space \mathbb{R}^d and densities given by

$$(5) \quad p_{\boldsymbol{\theta}}(\mathbf{x}) = \exp\{(\mathbf{A}^\top \mathbf{x}, \boldsymbol{\theta}) - \psi(\boldsymbol{\theta})\}, \quad \mathbf{x} \in S(\mathbf{V}), \boldsymbol{\theta} \in \mathbb{R}^d,$$

³This particular choice of the dominating measure will lead to Poisson and product multinomial likelihoods. More generally, much of our analysis carries over with other choices of dominating measure, for example, the ones for which conditions (A1)–(A4) in Rinaldo, Fienberg and Zhou (2009) hold.

where $\psi(\boldsymbol{\theta}) = \log(\int_{S(V)} \exp\{(\mathbf{A}^\top \mathbf{x}, \boldsymbol{\theta})\} d\nu_V(\mathbf{x}))$. This exponential representation is not the most parsimonious from the viewpoint of sufficiency. Indeed, let $\mathcal{T} = \{\mathbf{t} \in \mathbb{R}^d : \mathbf{t} = \mathbf{A}^\top \mathbf{x}, \mathbf{x} \in S(V)\}$ be the image of \mathbf{A} and $\mu_V = \nu_V \mathbf{A}^{-1}$ be the measure induced by \mathbf{A} . Then, by standard arguments [see, e.g., Brown (1986)], the distributions of the sufficient statistics $\mathbf{t} = \mathbf{A}^\top \mathbf{n}$ also form an exponential family, with density with respect to the base measure μ_V given by

$$(6) \quad q_{\boldsymbol{\theta}}(\mathbf{t}) = \exp((\mathbf{t}, \boldsymbol{\theta}) - \psi(\boldsymbol{\theta})), \quad \mathbf{t} \in \mathcal{T}, \boldsymbol{\theta} \in \mathbb{R}^d,$$

the same the log-partition function ψ and natural parameter space as in the original family.

It is now easy to see that the exponential family parametrization and the log-linear parametrization are equivalent. Indeed, for any \mathbf{n} and \mathbf{t} such that $\mathbf{t} = \mathbf{A}^\top \mathbf{n}$, and for any $\boldsymbol{\theta} \in \mathbb{R}^d$, the identity

$$(\mathbf{t}, \boldsymbol{\theta}) = (\mathbf{A}^\top \mathbf{n}, \boldsymbol{\theta}) = (\mathbf{n}, \mathbf{A}\boldsymbol{\theta}) = (\mathbf{n}, \boldsymbol{\mu}),$$

where $\boldsymbol{\mu} = \mathbf{A}\boldsymbol{\theta} \in \mathcal{M}$, implies these models can be equivalently parametrized by the linear subspace \mathcal{M} . If \mathbf{A} is of full rank, then the map $\boldsymbol{\theta} \mapsto \mathbf{A}\boldsymbol{\theta}$ is an isomorphism between \mathbb{R}^d and \mathcal{M} , while if $d > \dim(\mathcal{M})$, the natural parametrization is redundant and, in fact, nonidentifiable.

Throughout this article, will impose the following assumptions. Let \mathbf{V} be the matrix specifying the conditional Poisson sampling scheme.

- (A0) *Nontriviality*: the set $S(V)$ is nonempty.
- (A1) *Exhaustive sampling condition*: there does not exist any vector $\boldsymbol{\gamma} \in \mathcal{N}^\perp \setminus \{0\}$, such that $(\boldsymbol{\gamma}, \mathbf{n})$ is constant almost everywhere with respect to ν_V . In particular, for no cell $i \in \mathcal{I}$, $\mathbf{n}(i) = 0$, almost everywhere ν_V .
- (A2) *Integrity assumption*: $\{\mathbf{x} \in \mathbb{R}_{\geq 0}^T : \mathbf{V}\mathbf{x} = \mathbf{1}\} = \text{conv}(S(V))$.

Assumption (A1) guarantees that no linear constraints hold, other than the ones specified by \mathcal{N} , and it prevents the sampling constraints from introducing structural zeros. Even though we can easily extend our analysis to deal with structural zeros, we do not provide the details here. Assumption (A2) is technical, and it is used in Theorem 3 below to unify the conditions for existence of the MLE across different sampling schemes. If (A2) is not in effect, checking for existence of the MLE can become computationally infeasible, depending on \mathbf{V} . The Poisson, product multinomial and Poisson–multinomial schemes automatically satisfy (A2).

2.1. The effects of sampling constraints. We conclude this section by studying the effect of the sampling constraints on the estimability of the natural and log-linear parameters. We show that imposing linear sampling restrictions results in nonidentifiability of the corresponding natural exponential family (5), to the extent that only certain linear combinations of the

natural parameters, which depend only on the subspace \mathcal{N} , are estimable. For the log-linear parameters, only $\Pi_{\mathcal{M} \ominus \mathcal{N}} \boldsymbol{\mu}$ is estimable, which implies that the number of estimable parameters is $\dim(\mathcal{M} \ominus \mathcal{N}) = d - m$.

We define the following equivalence relation on \mathbb{R}^d : for $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$, $\boldsymbol{\theta}_1 \stackrel{\mathcal{N}}{\sim} \boldsymbol{\theta}_2$ if and only if $\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2 \in \mathcal{Z}$, where

$$(7) \quad \mathcal{Z} := \{\boldsymbol{\zeta} \in \mathbb{R}^d : A\boldsymbol{\zeta} \in \mathcal{N}\}.$$

For any $\boldsymbol{\theta} \in \mathbb{R}^d$, we then write $\boldsymbol{\theta}_{\mathcal{N}} := \{\boldsymbol{\theta}^* : \boldsymbol{\theta} \stackrel{\mathcal{N}}{\sim} \boldsymbol{\theta}^*\}$ for the equivalence class containing $\boldsymbol{\theta}$, and $\Theta_{\mathcal{N}} := \{\boldsymbol{\theta}_{\mathcal{N}}, \boldsymbol{\theta} \in \mathbb{R}^d\}$ for the set of equivalent classes corresponding to the equivalence relation $\stackrel{\mathcal{N}}{\sim}$. For simplicity, below we assume that the matrices A and V are of full rank, but the same conclusions hold with d replaced by $\text{rank}(A)$.

LEMMA 1. *Consider the exponential family (5), with A of full rank d , and suppose that conditions (A0) and (A1) hold.*

- (i) *The set $\Theta_{\mathcal{N}}$ is a vector space of dimension $d - m$ isomorphic to $\mathcal{M} \ominus \mathcal{N}$, and is comprised of parallel m -dimensional affine subspaces of \mathbb{R}^d .*
- (ii) *The family is nonidentifiable: any two points $\boldsymbol{\theta}_1 \stackrel{\mathcal{N}}{\sim} \boldsymbol{\theta}_2$ specify the same distribution. In fact, this family is parametrized by $\Theta_{\mathcal{N}}$, or, equivalently, by $\mathcal{M} \ominus \mathcal{N}$. Therefore, it is of order $d - m$.*

Using standard minimality arguments, nonidentifiability of the natural parameters can be easily resolved by redefining a smaller exponential family of order $d - m$ using as a new design matrix any full-rank matrix whose column span is $\mathcal{M} \ominus \mathcal{N}$; for this fully-identifiable family, the natural parameter space is \mathbb{R}^{d-m} . Concretely, we assume, without loss of generality, that the matrix A is of the form

$$(8) \quad A = (B \ V),$$

where V is the $I \times m$ matrix of sampling restrictions whose rows span \mathcal{N} and B is a $I \times (d - m)$ matrix whose row space is $\mathcal{M} \ominus \mathcal{N}$. Then, replacing A with B in (5) will produce a full and minimal exponential family.

To illustrate this point, we show that the log-likelihood function (3) for the product multinomial sampling scheme can be more conveniently parametrized by $\mathcal{M} \ominus \mathcal{N}$ instead of the nonconvex set $\tilde{\mathcal{M}}$. For any $\boldsymbol{\beta} \in \mathcal{M} \ominus \mathcal{N}$, let

$$(9) \quad \ell^M(\boldsymbol{\beta}) := (\mathbf{n}, \boldsymbol{\beta}) - \sum_{j=1}^m N_j \log(\exp(\boldsymbol{\beta}), \boldsymbol{\chi}_j) - \sum_{i \in \mathcal{I}} \log \mathbf{n}(i)!.$$

LEMMA 2. *The sets $\mathcal{M} \ominus \mathcal{N}$ are $\tilde{\mathcal{M}}$ homeomorphic and, for each pair of homeomorphic vectors $\boldsymbol{\mu} \in \tilde{\mathcal{M}}$ and $\boldsymbol{\beta} \in \mathcal{M} \ominus \mathcal{N}$, $\tilde{\ell}_{\mathcal{L}}(\boldsymbol{\mu}) = \ell^M(\boldsymbol{\beta})$.*

The form of the likelihood in (9) is better suited for computations, as we show in Fienberg and Rinaldo (2012).

Under the conditions of Lemma 1, the Fisher information matrix at θ has rank $d - m$, for each $\theta \in \mathbb{R}^d$. To see this, notice that the Fisher information matrix $I(\theta)$ at θ is $\text{Cov}_\theta(A^\top \mathbf{n})$, where Cov_θ denotes the covariance operator evaluated using the distribution parametrized by θ . Then, for any ζ in the set \mathcal{Z} defined in (7), the linear form $(A^\top \mathbf{n}, \zeta)$ is constant almost everywhere and therefore has zero variance. This is equivalent to $\zeta^\top I(\theta) \zeta = 0$, so that $\text{rank}(I(\theta)) = \dim(\mathcal{Z}^\perp) = d - m$, for all θ .

3. Theory of maximum likelihood estimation. We now provide a systematic treatment of maximum likelihood estimation for the natural and log-linear parameters, within the framework of the theory of discrete extended exponential families with linear sufficient statistics. We refer the reader to Barndorff-Nielsen (1978) and Brown (1986) for classic references and Csizsár and Matúš (2001, 2003, 2005, 2008) for advanced treatments. In our setting, Geyer (2009) and Rinaldo, Fienberg and Zhou (2009) are particularly relevant. For the reader's convenience, we briefly review the aspects of this theory that are relevant to our problem in Appendix A.

3.1. Existence of the MLE. We prove a general necessary and sufficient condition for existence of the MLE that applies to any conditional Poisson sampling scheme satisfying assumptions (A0)–(A2). Unlike existing results, these conditions directly translate into usable algorithms for checking for the existence of the MLE, as described in Fienberg and Rinaldo (2012).

For any design matrix A , we denote by $C_A := \text{cone}(A^\top)$ the polyhedral cone spanned by the rows of A . Following Eriksson et al. (2006), we call C_A the *marginal cone* of A .

THEOREM 3. *Assume conditions (A0)–(A2) and let A be any matrix with column span \mathcal{M} . The MLE of θ_N (or, equivalently, of $\Pi_{\mathcal{M} \ominus N} \mu$) exists and is unique if and only if $\mathbf{t} = A^\top \mathbf{n} \in \text{ri}(C_A)$.*

This result is a nontrivial application of a well-known result about existence of MLE in exponential families (viz., Theorem 13 in Appendix A), and it subsumes previous results of Haberman (1974) and Eriksson et al. (2006), because it provides a unified condition that applies to all conditional Poisson sampling schemes satisfying the integrality assumption (A2). To see how Theorem 3 differs from Theorem 13, a direct application of the latter yields that the MLE exists if and only if \mathbf{t} belongs to the interior of the $(d - m)$ -dimensional polyhedron

$$C_V := \text{conv}(\{\mathbf{t} : \mathbf{t} = A^\top \mathbf{x}, \mathbf{x} \in \mathbb{N}^Z, V^\top \mathbf{x} = \mathbf{1}\}).$$

For Poisson sampling, this polyhedron is in fact the marginal cone, and, for multinomial sampling, it is the polytope $\{V\mathbf{x} : \mathbf{x} \in \text{conv}(A)\}$. Under prod-

uct multinomial sampling, C_V is the Minkowski addition [see, e.g., Ziegler (1995), Schrijver (1998)] of m polytopes, one for each multinomial, while under Poisson–multinomial scheme it is the Minkowski sum of a polyhedral cone and as many polytopes as multinomial constraints. Even though it has smaller ambient dimension than the marginal cone, C_V is a geometric object that can be rather difficult to handle, both computationally and theoretically. In contrast, we show that, for any sampling scheme satisfying conditions (A0)–(A2), it is in fact sufficient to deal with the polyhedral cone C_A , which is simpler to describe and analyze, both algorithmically and in theory; see the supplementary material Fienberg and Rinaldo (2012). In Rinaldo, Petrović and Fienberg (2011) we provide various examples of how Theorem 3 can be used to simplify the task of characterizing existence of the MLE for otherwise complicated models for networks and random graphs. These particular models are based on product multinomial sampling constraints, in which case Theorem 3 yields what is known in polyhedral geometry as the Cayley trick.

3.2. Parameter estimability. We now turn to the issue of estimability of the natural and log-linear parameters when the MLE does not exist. In our analysis, we rely on the key notion of facial sets, originally introduced in a slightly different form by Geiger, Meek and Sturmfels (2006).

DEFINITION 4. For a log-linear subspace \mathcal{M} , a set $\mathcal{F} \subseteq \mathcal{I}$ is a facial set of \mathcal{M} , when, for some $\boldsymbol{\mu} \in \mathcal{M}$,

$$\begin{aligned}
 \boldsymbol{\mu}(i) &= 0 & \text{if } i \in \mathcal{F}, \\
 \boldsymbol{\mu}(i) &< 0 & \text{if } i \notin \mathcal{F}.
 \end{aligned}$$

Equivalently, \mathcal{F} is a facial set of \mathcal{M} when, for any design matrix A for \mathcal{M} (not necessarily of full column rank), there exists some $\mathbf{c} \in \mathbb{R}^d$ such that

$$\begin{aligned}
 (\mathbf{a}_i, \mathbf{c}) &= 0 & \text{if } i \in \mathcal{F}, \\
 (\mathbf{a}_i, \mathbf{c}) &< 0 & \text{if } i \notin \mathcal{F},
 \end{aligned} \tag{10}$$

where \mathbf{a}_i denotes the i th row of A . Facial sets encode combinatorial and geometric properties of the log-linear subspace \mathcal{M} which turn out to be crucial to our analysis. We summarize these properties in the next lemma.

LEMMA 5. *Let A be a design matrix of \mathcal{M} . The lattice of facial sets of \mathcal{M} is isomorphic to the face lattice of the marginal cone C_A . In particular, \mathcal{F} is a facial set of \mathcal{M} if and only if $\{\mathbf{a}_i, i \in \mathcal{F}\}$ span the face of C_A isomorphic to \mathcal{F} .*

Using this result, we can paraphrase Theorem 3 as follows [compare with Theorem 3.2 in Haberman (1974)]:

COROLLARY 6. *The MLE exists if and only if $\text{supp}(\mathbf{n})$ is not contained in any facial set of \mathcal{M} .*

We describe algorithms for determining facial sets and for using the previous corollary in Fienberg and Rinaldo (2012).

3.2.1. *Estimability of the natural parameters.* In this section, we rely on arguments proposed in Rinaldo, Fienberg and Zhou (2009) to study the estimability of the natural parameters. Let C_V denote the convex support of the family arising from a conditional Poisson scheme specified by a constraint matrix V ; see Appendix A. Suppose that the observed sufficient statistics $\mathbf{t} = A^\top \mathbf{n}$ belong to the relative interior of face F_V of C_V of dimension d_F . Thus, the MLE of the natural parameters for the original family, supported on $S(V)$, is nonexistent, but the MLE of the natural parameter of the extended family supported F_V is well defined. Theorem 7 below generalizes Lemma 1 by showing that, when the MLE does not exist, the linear combinations of the natural parameters that are estimable are determined, not only by the deterministic linear subspace arising from the sampling constraints, but also by the random linear subspace spanned by the normal cone to the face F of the marginal cone C_A containing $A^\top \mathbf{n}$ in its relative interior. As for the log-linear parameter, nonexistence of the MLE entails that only points in $\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})$ are estimable, where \mathcal{F} is the random facial set corresponding to F .

In preparation for the result, we need to set up some additional notation. By Lemma 15 in Appendix B, there exists one face F of C_A of dimension $m + d_F$ that contains F_V , with facial set \mathcal{F} . Let N_F be the normal cone to F and $\mathcal{L}_F \subset \mathbb{R}^d$ be the linear subspace spanned by N_F , so that $\dim(\mathcal{L}_F) = d - m - d_F$ (recall that, without loss of generality, we assume C_A to be full-dimensional). We further define the linear subspace

$$\mathcal{N}_F := \{A\beta, \beta \in \mathcal{Z} + \mathcal{L}_F\},$$

where \mathcal{Z} is given in (7). Just like in Lemma 1, we define the following equivalence relation on \mathbb{R}^d : $\theta_1 \sim_{\mathcal{N}_F} \theta_2$ if and only if $\theta_1 - \theta_2 \in \mathcal{Z} + \mathcal{L}_F$, and write $\theta_{\mathcal{N}_F}$ for the equivalence class containing θ . Finally, $\Theta_{\mathcal{N}_F} := \{\theta_{\mathcal{N}_F}, \theta \in \mathbb{R}^d\}$.

THEOREM 7. *Consider the exponential family (5), with A of full rank d , and suppose that conditions (A0)–(A2) hold. Let F_V be a face of the convex support and \mathcal{F} the corresponding facial set of the normal cone.*

- (i) *For any $\theta \in \mathbb{R}^d$, the set $\theta_{\mathcal{N}_F}$ is an affine subspace of \mathbb{R}^d of dimension $m + \dim(\mathcal{L}_F) = d - d_F$. The set $\Theta_{\mathcal{N}_F}$ is a d_F -dimensional dimensional vector space isomorphic to $\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})$ and is comprised of parallel $(d - d_F)$ -dimensional affine subspaces of \mathbb{R}^d .*

- (ii) *The extended family corresponding to F_V is non-identifiable: any two points $\theta_1 \stackrel{\mathcal{N}_F}{\sim} \theta_2$ specify the same distribution. In fact, the family is parametrized by $\Theta_{\mathcal{N}_F}$, or, equivalently, by $\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})$. Therefore, it is of order d_F .*

The main point of Theorem 7 is that only natural parameters in $\Theta_{\mathcal{N}_F}$ [or the log-linear parameters in $\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})$] are estimable, with both sets being now random. In principle, nonidentifiability of the natural parameters, due to a nonexistent MLE, can be resolved using the same procedure of reduction to minimality described in the remarks following Lemma 1: identify a set of linearly independent vectors in \mathbb{R}^I spanning $\mathcal{M} \cap \mathcal{N}_F^\perp$, and use them to build a new design matrix of dimension $I \times d_F$. However, unlike the reduction to minimality carried out to remove the effect of the sampling constraints, which is design-dependent but not data-dependent, this reduction depends on the random subspace \mathcal{N}_F (the randomness arising from the exposed face F). Furthermore, while the sampling constraint reduction is easy to implement, since the matrix V is known, this second reduction requires us to compute a basis for \mathcal{L}_F , the linear space spanned by the normal cone to F . For the mean value parameter, the problem is to compute the facial set associated to the face F based solely on the observed sufficient statistics \mathbf{t} , which amounts to identifying the face of C_A containing \mathbf{t} in its relative interior. In general, both of these tasks are highly nontrivial, due to the combinatorial complexity of the face lattice of C_A ; see the examples in Section 4. In the supplementary material [Fienberg and Rinaldo (2012)], we describe algorithms for accomplishing these tasks.

As a corollary to Theorem 7, we can obtain each family in the extended family via a conditional Poisson sampling scheme that forces the base measure to be supported on F_V , or equivalently, by requiring that the cells in \mathcal{F}^c have zero probability of containing positive counts. In this case, it is clear that assumption (A1) is violated. As a result, we can view each such family as a log-linear model under Poisson sampling scheme containing structural zeros along the (random) coordinates \mathcal{F}^c . This is in fact consistent with the interpretation by Barndorff-Nielsen (1978), page 156, of the extended MLE as a conditional MLE, given that sufficient statistics lie on the boundary of the convex support. We formalize this observation in the next result.

COROLLARY 8. *Each face F of C_V of dimension $0 \leq d_F \leq d - m$ can be obtained as the convex support corresponding to the conditional Poisson scheme with constraint subspace \mathcal{N}_F , where $\dim(\mathcal{N}_F) = d - d_F$.*

Using the same arguments as in the remarks following Lemma 1, we also see that the Fisher information matrix at the extended MLE has rank $d_F < d$, and therefore, is rank-deficient. This remains the case, even after accounting for the sampling constraints. Statistically, the singularity of the

observed Fisher information implies that the standard errors are not defined. From an algorithmic standpoint, this observation implies that the Newton–Raphson method for computing the MLE is bound to run into numerical instabilities, due to the fact that the Hessian matrix of the log-likelihood function is singular at any optimum [an issue illustrated empirically in Fienberg and Rinaldo (2007)]. Furthermore, Corollary 2.8 in Rinaldo, Fienberg and Zhou (2009) shows that, under a nonexistent MLE, every point in the normal cone N_F to the face F containing the observed sufficient statistics is a (random) direction of recession of the negative log-likelihood function, so that there are infinitely many directions of maximal increase of the log-likelihood function.

3.2.2. Estimability of the mean value parameters under Poisson and product multinomial schemes. We now specialize our analysis to the case of Poisson and product multinomial sampling schemes. Besides their popularity, the main reason for focusing on these two particular sampling schemes is that the estimates of the cell mean values are highly interpretable. Under the Poisson scheme, the cell mean values are just the expected cell counts, while under the product multinomial scheme they are the conditional expectations of the cell counts given the grand total (in the multinomial case) or given the total counts in the portions of the table associated with the partitions used to define the product multinomial constraints. For other conditional Poisson sampling schemes, not only are the conditional cell mean values difficult to compute due to the unknown normalizing constant, but they are also less interpretable.

Following Lauritzen (1996), we consider $\overline{M} = \text{cl}(\{\exp(\boldsymbol{\mu}), \boldsymbol{\mu} \in \mathcal{M}\})$, the closure of the set of all cell mean values for a log-linear subspace \mathcal{M} . Thus, $\mathbf{m} \in \overline{M}$ if and only if $\mathbf{m} = \lim_n e^{\boldsymbol{\mu}_n}$, for some sequence $\{\boldsymbol{\mu}_n\}_n \subset \mathcal{M}$. Lauritzen (1996) calls the set \overline{M} the *extended log-affine model*.

THEOREM 9. *Let \mathbf{t} be the observed sufficient statistics, and let \mathcal{F} be facial set corresponding to the face of C_A containing \mathbf{t} in its relative interior. The MLE of the cell mean vector exists, is unique and identical under Poisson and product multinomial if and only if $\mathcal{F} = \mathcal{I}$. If $\mathcal{F} \subsetneq \mathcal{I}$, there exists one point $\widehat{\mathbf{m}}^e$ in \overline{M} such that $\widehat{\mathbf{m}}^e = \lim_n \exp(\boldsymbol{\mu}_n)$, where $\{\boldsymbol{\mu}_n\}_n \subset \mathcal{M}$ is any optimizing sequence such that*

$$\lim_n \ell^P(\boldsymbol{\mu}_n) = \sup_{\boldsymbol{\mu} \in \mathcal{M}} \ell^P(\boldsymbol{\mu}) \quad \text{and} \quad \lim_n \ell^M(\boldsymbol{\mu}_n) = \sup_{\tilde{\boldsymbol{\mu}} \in \widetilde{\mathcal{M}}} \ell^M(\tilde{\boldsymbol{\mu}}).$$

Furthermore, $\text{supp}(\widehat{\mathbf{m}}^e) = \mathcal{F}$ and $\Pi_{\mathcal{M}} \mathbf{n} = \Pi_{\mathcal{M}} \widehat{\mathbf{m}}^e$.

This result shows that, for any observed table \mathbf{n} , the log-likelihood functions in both sampling schemes admits always a unique maximizer, $\widehat{\mathbf{m}}^e$. Though supported only on the facial set associated with \mathbf{t} , this vector exhibits exactly the same features as the “ordinary” MLE: it is the unique

point $\widehat{\mathbf{m}}^e \in \overline{M}$ such that $A^\top \widehat{\mathbf{m}}^e = A^\top \mathbf{n}$ and provided that $\mathcal{N} \subset \mathcal{M}$, maximizes both the Poisson and product multinomial likelihoods. The substantial difference is that $\widehat{\mathbf{m}}^e$ has positive coordinates only along the cells in the facial set \mathcal{F} . Theorem 9 generalizes Theorem 4.8 in Lauritzen (1996). The improvement consists of identifying exactly the supports of the limit points in \overline{M} , which are precisely the facial sets of C_A .

DEFINITION 10. The vector $\widehat{\mathbf{m}}^e$ is the *extended MLE* of \mathbf{m} and the zeros appearing in along the coordinates in $\mathcal{F}^c = \mathcal{I} \setminus \mathcal{F}$ are called the *likelihood zeros*.

The term likelihood zeros highlight the fact that those zero counts, though arising as sampling and not as structural zeros, have a significant impact on the likelihood function and its optimizers.

3.3. *The geometry of the extended Poisson family.* The results of Theorem 9 suggest that, for the Poisson and product multinomial schemes, we could, in fact, take the set \overline{M} to be the cell mean value parameter space for the extended exponential family of distributions for the actual contingency table, not its sufficient statistics. We formalize this idea by relying on geometric considerations. For ease of readability, and without loss of generality, we focus on the Poisson sampling scheme, and only sketch how our results apply also to product multinomial cases.

For a vector $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$, let

$$\mathbf{u}^+ = \{\max\{\mathbf{u}(i), 0\}, i \in \mathcal{I}\} \quad \text{and} \quad \mathbf{u}^- = \{\min\{\mathbf{u}(i), 0\}, i \in \mathcal{I}\},$$

so that $\mathbf{u} = \mathbf{u}^+ - \mathbf{u}^-$ and $\text{supp}(\mathbf{u}^+) \cap \text{supp}(\mathbf{u}^-) = \emptyset$. Furthermore, for any pair of nonnegative vectors \mathbf{x} and \mathbf{u} in $\mathbb{R}^{\mathcal{I}}$, write

$$\mathbf{x}^{\mathbf{u}} = \prod_i \mathbf{x}(i)^{\mathbf{u}(i)}$$

for the associated monomial. Following Geiger, Meek and Sturmfels (2006), page 1469 and Lemma A.1, we consider the toric variety $X_{\mathcal{M}}$ corresponding to the log-linear model \mathcal{M} .

DEFINITION 11. The nonnegative toric variety $X_{\mathcal{M}}$ associated to the log-linear subspace \mathcal{M} is the set of all vectors $\mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}}$ such that

$$(11) \quad \mathbf{x}^{\mathbf{u}^+} = \mathbf{x}^{\mathbf{u}^-} \quad \forall \mathbf{u} \in \mathcal{M}^\perp.$$

Geometrically, $X_{\mathcal{M}}$ is the intersection of the solution set of a system of polynomial equations with the nonnegative orthant. It is easy to see that any $\mathbf{m} > 0$ such that $\log(\mathbf{m}) \in \mathcal{M}$ satisfies (11). Equation (11) can still hold, however, when some of the coordinates of \mathbf{m} are zero. Finally, for any $\boldsymbol{\xi} \in C_A$, consider the polyhedron

$$(12) \quad P_{\boldsymbol{\xi}} = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}} : A\mathbf{x} = \boldsymbol{\xi}\}.$$

For a given sufficient statistic $\mathbf{t} = \mathbf{A}\mathbf{n}$, the set of lattice points in $P_{\mathbf{t}}$, known as the *fiber* of \mathbf{t} , consists of all possible tables having the same sufficient statistics as the observed table \mathbf{n} .

- THEOREM 12. (i) $\overline{M} = X_{\mathcal{M}}$.
(ii) For any nonzero $\mathbf{m} \in X_{\mathcal{M}}$, $\text{supp}(\mathbf{m})$ is a facial set of C_A .
(iii) The linear map $A: \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}^d$, given by $\mathbf{m} \mapsto \mathbf{A}\mathbf{m}$, defines a homeomorphism between $X_{\mathcal{M}}$ and C_A .
(iv) For any observable sufficient statistic $\mathbf{t} = \mathbf{A}\mathbf{n}$, $\{\hat{\mathbf{m}}^e\} = X_{\mathcal{M}} \cap P_{\mathbf{t}}$ and $\hat{\mathbf{m}}^e \in \text{ri}(P_{\mathbf{t}})$.

Part (i) of Theorem 12 is due to Geiger, Meek and Sturmfels (2006), while a slightly less general version of part (iii) is a standard result in the algebraic statistics literature; see, for example, Pachter and Sturmfels (2005), Drton, Sturmfels and Sullivant (2009).

Overall, Theorem 12 shows that the set \overline{M} is homeomorphic to the marginal cone C_A and, therefore, as anticipated, we can use it as a legitimate mean value space for the extended family of the cell counts. The advantage of \overline{M} over C_A is its direct interpretability in terms of cell mean values. This result extends directly to the multinomial sampling scheme. In this case, A specifies a homeomorphism between $\{\mathbf{x} \in X_{\mathcal{M}}: \sum_i \mathbf{x}_i = 1\}$ and $P_A = \text{conv}(A)$, which is known in algebraic geometry as the moment map; see Fulton (1993), Ewald (1996). In fact, under multinomial scheme, the extended mean-value space can be taken to be the intersection of $X_{\mathcal{M}}$ with the probability simplex in $\mathbb{R}^{\mathcal{I}}$. Furthermore, since \mathcal{M} contains the constant functions, P_A and C_A have identical facial sets. For product multinomial sampling schemes, a characterization of the mean value space analogous to the one given in Theorem 12 is also possible, though somewhat more involved. We refer the reader to Morton (2008) for details and a different derivation. In this particular case, the convex support arises as a Minkowski sum of polytopes, one for every multinomial. Then, the proof of Theorem 3 reveals that facial sets of the convex support are also facial sets of the marginal cone, even though the opposite is not true. See Rinaldo, Petrović and Fienberg (2011) for an application of these results to network models.

Finally, part (iv) of Theorem 12 shows that the extended MLE is the only point in $P_{\mathbf{t}}$ satisfying the log-linear model conditions. This result can be also interpreted in terms of I-divergence projections [Csiszár (1975, 1989)], and provides the geometric basis for showing convergence of iterative methods for extended maximum likelihood estimation such as the iterative proportional scaling algorithm of Darroch and Ratcliff (1972). In the interest of space, we do not pursue this analysis.

4. Inference under a nonexistent MLE. We have shown that when the MLE does not exist, only some of the model parameters (both under the natural and mean-value parametrization) are estimable, and we have identified

the parameters that can instead be estimated within the extended family. Thus, when the MLE is nonexistent, statistical inference is still feasible, but only for the reduced family whose parameters are fully estimable.

As described at the end of Section 3.2.1, we can obtain the relevant extended exponential family by computing a new random design matrix $A_{\mathcal{F}}$ whose column span is $\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})$, where \mathcal{F} is the random facial set corresponding to the face F of the marginal cone containing the sufficient statistics in its relative interior. We can then use this new design matrix to specify a new exponential family as in (5), where *only the cells \mathcal{F} have positive probability of being observed*. We carry out inference within this extended family or, equivalently, conditionally on the sufficient statistics being on the face F , as advocated by Barndorff-Nielsen (1978), page 156. By Corollary 8, this is equivalent to treating the coordinates in \mathcal{F} as if they were structural zeros. Thus, dealing with a nonexistent MLE reduces, in practice, to fitting the same log-linear model under the additional (random) constraints that the cells in \mathcal{F}^c , which are not estimable, be treated as structural zeros. The same approach is also advocated in Geyer (2009). In practice, this entails replacing the MLE with the extended MLE and, quite importantly, adjusting the number of degrees of freedom, now to be computed as the difference between the cardinality of the facial set $|\mathcal{F}|$ (i.e., the number of cell mean values that can be estimated), and the number of estimable parameters, namely $\dim(\pi_{\mathcal{F}}(\mathcal{M} \ominus \mathcal{N})) = \dim(F) - m$. Using the adjusted number of degrees of freedom, asymptotic χ^2 tests for goodness of fit [see, e.g., Read and Cressie (1988)] can then still be applied. Algorithms for carrying out the numerical tasks just described are presented in the supplementary material [Fienberg and Rinaldo (2012)].

5. Examples of likelihood zeros. Below, we illustrate by means of examples various practical aspects of goodness-of-fit testing when the MLE is nonexistent, and we show how to appropriately adjust the number of degrees of freedom. We will focus on hierarchical log-linear models [see, e.g., Bishop, Fienberg and Holland (1975)], and refer the reader to Dobra et al. (2009) and Rinaldo, Petrović and Fienberg (2011) for other examples of this kind.

Our polyhedral characterization of the conditions for the existence of the MLE permits to generate novel examples of patterns of sampling zeros causing nonexistence of the MLE for hierarchical log-linear models without producing null margins, an instance that is virtually ignored in all statistical software. As pointed out by Fienberg and Rinaldo (2007), the R [R Development Core Team (2005)] routines `loglin` and `glm`, as well as virtually any other software for inference and model selection for log-linear models, does not detect nonexistence and report the unadjusted, incorrect, numbers of degrees of freedom for all the examples below. In the analysis of sparse tables, it is also common practice to add small positive quantities to the zero

cells, in order to avoid numerical issues with the computation of the MLE. We remain highly skeptical of the numerical advantages of this ad-hoc procedure, and remark that such adjustments will make it impossible to detect nonexistence of the MLE and to distinguish the estimable parameters.

The examples of likelihood zeros in Examples 2–4 suggest that the combinatorial complexity of hierarchical log-linear models, measured by the number of facets of the marginal cone, can be quite significant. In the reported examples, as well as in many other experiments we conducted, for many models the number of facets associated with zero margins appears to be much smaller than the total number of facets, indicating that, at least combinatorially, likelihood zeros associated to positive margins are much more frequent (though never detected). Below we use the classic notation to represent the generating class of a hierarchical log-linear model; for example, see Bishop, Fienberg and Holland (1975). Empty cells indicate positive counts. All the calculations were carried out in `polymake` [Gawrilow and Joswig (2000)].

EXAMPLE 1. The 2^3 table and the model $[12][13][23]$ of no-second-order interaction. The MLE is not defined because the two likelihood zeros expose one of the 16 facets of the marginal cone. This example, due to Haberman (1974), was the only published example a log-linear model with nonexistent MLE and positive margins; see Fienberg and Rinaldo (2007), Section 5, for a general result concerning binary K -way tables and the model of no- $(K - 1)$ st interaction.

| | | | | | |
|---|--|--|--|--|--|
| 0 | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

The dimension of the log-linear subspace for this model, or, equivalently, of the marginal cone, is 7, leaving 1 degree of freedom when the MLE exists. However, because of the likelihood zeros, inference can only be made for the 6-dimensional exposed facet. Since the cardinality of the associated facial set \mathcal{F} is also 6, the resulting extended log-linear model is the saturated model on \mathcal{F} .

EXAMPLE 2. The 3^3 table and the model $[12][13][23]$. The MLE is not defined because the pattern of likelihood zeros exposes one of the 207 facets of the marginal cone. Of all the facets, only 27 are associated to zero margins.

| | | | | | | | | |
|---|---|--|--|---|---|---|--|---|
| 0 | | | | | | 0 | | 0 |
| | | | | | | | | 0 |
| | | | | 0 | 0 | | | 0 |
| 0 | 0 | | | 0 | | | | |
| | | | | | | | | |
| | | | | | | | | |

The dimension of the facet is 18, which is also the cardinality of the facial set for this configuration of likelihood zeros. As in the previous example, this defines the saturated model on \mathcal{F} , giving 0 adjusted degrees of freedom and making χ^2 approximations not applicable.

Under the same log-linear model, the MLE does not exist also when the following pattern of zeros arises:

| | | |
|---|--|--|
| 0 | | |
| 0 | | |
| | | |

| | | |
|--|---|---|
| | | |
| | | 0 |
| | 0 | 0 |

| | | |
|---|--|---|
| 0 | | 0 |
| 0 | | |
| | | |

In this example, the zeros displayed in bold are not likelihood zeros, but the others are. Indeed, their presence or absence has no effect on the existence of the MLE. Furthermore, when the extended MLE is computed, the boldfaced zero counts will be replaced by positive entries, while the likelihood zeros will stay zero. The number of degrees of freedom in this example is 3, because the total number of estimable cell mean values is 21, and the number of parameters for the reduced model is 18.

In our last example, the MLE is defined, despite the table being very sparse, because no facet of the marginal cone is exposed [source: Fienberg and Rinaldo (2007)].

| | | |
|---|---|---|
| | 0 | 0 |
| 0 | | 0 |
| 0 | 0 | |

| | | |
|---|---|---|
| 0 | 0 | |
| | 0 | 0 |
| 0 | | 0 |

| | | |
|---|---|---|
| 0 | | 0 |
| 0 | 0 | |
| | 0 | 0 |

EXAMPLE 3. The $4 \times 4 \times 4$ table and the model [12][13][23]. The MLE is not defined because the pattern of zeros exposes one of 113,740 facets of the marginal cone [source: Eriksson et al. (2006)]. Of these, only 48 are associated to zero margins.

| | | | |
|---|---|---|--|
| 0 | 0 | 0 | |
| 0 | 0 | | |
| 0 | | | |
| | | | |

| | | | |
|--|---|---|---|
| | 0 | 0 | |
| | 0 | | |
| | | | |
| | 0 | 0 | 0 |

| | | | |
|---|--|---|---|
| | | 0 | |
| | | | |
| 0 | | 0 | 0 |
| | | 0 | 0 |

| | | | |
|---|---|--|---|
| | | | |
| 0 | 0 | | 0 |
| 0 | | | 0 |
| | | | 0 |

EXAMPLE 4. The 3^4 table and the 4-cycle model [12][14][23][34]. The MLE is not defined because the pattern of zeros exposes one of the 1116 facets of the marginal cone. Of these, only 36 are associated to zero margins.

| | | |
|---|--|---|
| 0 | | 0 |
| | | 0 |
| | | |

| | | |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 0 | 0 |
| 0 | 0 | |

| | | |
|---|---|---|
| 0 | 0 | 0 |
| | 0 | 0 |
| | 0 | |

| | | |
|---|---|---|
| 0 | | 0 |
| 0 | 0 | 0 |
| 0 | 0 | 0 |

| | | |
|---|---|---|
| 0 | | |
| 0 | 0 | 0 |
| 0 | 0 | |

| | | |
|--|---|---|
| | | |
| | 0 | 0 |
| | 0 | |

| | | |
|---|---|---|
| 0 | | 0 |
| | | 0 |
| 0 | 0 | 0 |

| | | |
|---|---|--|
| 0 | | |
| | | |
| 0 | 0 | |

| | | |
|---|---|---|
| 0 | 0 | 0 |
| | 0 | 0 |
| 0 | 0 | 0 |

6. Algorithms for extended maximum likelihood estimation. In the supplementary material [Fienberg and Rinaldo (2012)], we apply the theory developed in this article to develop efficient algorithms for extended maximum likelihood estimation in log-linear models under Poisson and product multinomial schemes [for which the key integrality assumption (A2) is satisfied] that are applicable to high-dimensional models and large tables. Some of these algorithms are implemented in a MATLAB toolbox available at <http://www.stat.cmu.edu/~arinaldo/ExtMLE/>. The final output of our procedure is the set of estimable mean value and natural parameters.

APPENDIX A: EXTENDED EXPONENTIAL FAMILIES

In this appendix we provide a brief review of the theory of extend families and its relevance for log-linear models. Along with classic references on exponential families [Barndorff-Nielsen (1978), Brown (1986), Čencov (1982), Letac (1992)] and generalizations by Csiszár and Matúš (2001, 2003, 2005, 2008), we refer the reader to Rinaldo, Fienberg and Zhou (2009) and Geyer (2009) for treatments more directly relevant to our problem.

Consider a log-linear model under conditional Poisson sampling scheme specified by a sampling matrix V of rank m and a design matrix A of the form (8), where B is of full-rank $d - m$. Then [see equation (6)], the distribution of the sufficient statistic $\mathbf{z} = B^\top \mathbf{n}$ form an exponential family of distributions \mathcal{E}_{C_V} on \mathbb{R}^{d-m} with densities

$$q_\theta(\mathbf{z}) = \exp((\mathbf{z}, \boldsymbol{\theta}) - \psi(\boldsymbol{\theta})), \quad \boldsymbol{\theta} \in \Theta,$$

with respect to the base measure $\mu_V = \nu_V^{-1}B$, and parameter space $\Theta = \mathbb{R}^{d-m}$. The *convex support* C_V of \mathcal{E}_{C_V} is the closure of the convex hull of the support of μ_V . In particular, P is a full-dimensional polyhedron in \mathbb{R}^{d-m} and, for every face F of C_V , F is the convex hull of some points in the support of μ_V . Given a realization \mathbf{z} of the sufficient statistics, the random set

$$(13) \quad \hat{\boldsymbol{\theta}}(\mathbf{z}) = \hat{\boldsymbol{\theta}} = \left\{ \boldsymbol{\theta}^* \in \Theta : q_{\boldsymbol{\theta}^*}(\mathbf{z}) = \sup_{\boldsymbol{\theta} \in \Theta} q_\theta(\mathbf{z}) \right\}$$

is the *maximum likelihood estimator* (MLE) of $\boldsymbol{\theta}$. If $\hat{\boldsymbol{\theta}} = \emptyset$ we say that the MLE does not exist. Existence of the MLE is fully characterized by the geometry of C_V , as the following well-known result indicates; see, for example, Theorem 5.5 in Brown (1986) or Theorem 9.13 in Barndorff-Nielsen (1978).

THEOREM 13. *For a minimal and full exponential family, the MLE $\hat{\boldsymbol{\theta}}$ exists and is unique if and only if $\mathbf{z} \in \text{ri}(P)$.*

Setting $\boldsymbol{\xi}(\boldsymbol{\theta}) = \int_{\mathbb{R}^{d-m}} \mathbf{z} q_\theta(\mathbf{z}) d\mu_V(\mathbf{z})$, because of the minimality of \mathcal{E}_{C_V} , one obtains the fundamental identity $\nabla \psi(\boldsymbol{\theta}) = \boldsymbol{\xi}(\boldsymbol{\theta}), \forall \boldsymbol{\theta} \in \Theta$, where ∇ indicates the gradient. In particular, if the MLE exists, it satisfies the equation $\hat{\boldsymbol{\theta}} = (\nabla \psi)^{-1}(\mathbf{z})$, which is equivalent to the moment equation $\boldsymbol{\xi}(\hat{\boldsymbol{\theta}}) = \mathbf{z}$.

For any proper face F of C_V , let μ_V^F be the restriction of μ_V to F . Then, μ_V^F determines a new exponential family of distributions, \mathcal{E}_F , with densities with respect to μ_V^F given by

$$q_{\boldsymbol{\theta}}^F(x) = \exp((\mathbf{z}, \boldsymbol{\theta}) - \psi^F(\boldsymbol{\theta})), \quad \boldsymbol{\theta} \in \Theta_F,$$

where the natural parameter space is $\Theta_F = \{\boldsymbol{\theta} \in \Theta : \exp(\psi^F(\boldsymbol{\theta})) < \infty\} = \Theta$, with $\psi^F(\boldsymbol{\theta}) = \log \int_{\mathbb{R}^{d-m}} \exp((\mathbf{z}, \boldsymbol{\theta})) d\mu_V^F(\mathbf{z})$. The convex support of this new family is F , and the existence result of Theorem 13 carries over: the MLE exists if and only if the observed sample \mathbf{z} belongs to $\text{ri}(F)$. However, since \mathcal{E}_F is supported on a lower-dimensional affine subspace of \mathbb{R}^{d-m} of dimension $d_F = \dim(F)$, it is no longer minimal, hence it is unidentifiable. Nonetheless, if $\mathbf{z} \in \text{ri}(F)$, the MLE of $\boldsymbol{\theta}$ is the set consisting of those $\boldsymbol{\theta}$ satisfying the first order optimality conditions

$$(14) \quad \mathbf{z} = \nabla \psi_F(\boldsymbol{\theta}) \quad \text{or, equivalently,} \quad \boldsymbol{\xi}^F(\boldsymbol{\theta}) = \mathbf{z},$$

where $\boldsymbol{\xi}^F(\boldsymbol{\theta}) = \int_{\mathbb{R}^{d-m}} \mathbf{z} q_{\boldsymbol{\theta}}^F(\mathbf{z}) d\mu_V^F(\mathbf{z})$.

The collection of distributions

$$\mathcal{E} = \bigcup_F \mathcal{E}_F$$

as F ranges over all the faces of C_V , including C_V itself, is known as the *extended exponential family* of distributions. With respect to such family \mathcal{E} , for any observed sample \mathbf{z} , the MLE, or *extended MLE*, is always well defined and is the set of solutions to (14), where F is the unique face containing \mathbf{z} in its relative interior.

APPENDIX B: PROOFS

This appendix contains the proofs of some results stated in the article. The remaining proofs can be found in the supplementary material Fienberg and Rinaldo (2012). Throughout, we assume familiarity with basic notions of polyhedral geometry; see Ziegler (1995), Schrijver (1998) and Rockafellar (1970) for in-depth treatments, and Section 2.1 of Rinaldo, Fienberg and Zhou (2009) for a brief review of the concepts directly relevant to our setting.

PROOF OF THEOREM 3. We first assume that A is of full rank d . If $\mathcal{N} = \{\mathbf{0}\}$, then the convex support is the d -dimensional polyhedral cone C_A , so the result follows directly from Theorem 13. Thus, throughout the remainder of the proof we consider the case $0 < \dim(\mathcal{N}) < d$. For now, we further assume that A is of the form (8).

By standard minimality arguments, we can work with the exponential family supported on $S = \{\mathbf{z} : \mathbf{z} = B^\top \mathbf{x}, \mathbf{x} \in \mathbb{N}^L, V^\top \mathbf{x} = 1\}$. By assumption (A1), the convex support C_V , which is the closure of the set $\text{conv}(S)$, is a full-dimensional polyhedron in \mathbb{R}^{d-m} . In particular, the parameter space is \mathbb{R}^{d-m} . The MLE exists and is unique if and only if $\mathbf{z} \in \text{ri}(C_V)$ by Theorem 13. We

now show that this happens if and only if $\mathbf{t} \in \text{ri}(C_A)$. We first use the integrality assumption (A2) to obtain a simpler representation of C_V .

LEMMA 14.

$$C_V = \{\mathbf{B}^\top \mathbf{x} : \mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}.$$

PROOF. Since $\{\mathbf{B}^\top \mathbf{x} : \mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}$ is a polyhedron (hence closed and convex), it must contain \bar{C}_V . To show the reverse inclusion, let $\mathbf{z}^* \in \{\mathbf{B}^\top \mathbf{x} : \mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}$. Then, $\mathbf{z}^* = \mathbf{B}^\top \mathbf{x}^*$ for some $\mathbf{x}^* \in \{\mathbf{x} : \mathbf{x} \in \mathbb{R}_{\geq 0}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}$. By the integrality assumption (A2),

$$\mathbf{x}^* \in \text{conv}(\{\mathbf{x} : \mathbf{x} \in \mathbb{N}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}),$$

which by linearity implies that $\mathbf{z}^* \in \text{conv}(\{\mathbf{B}^\top \mathbf{x} : \mathbf{x} \in \mathbb{N}^{\mathcal{I}}, \mathbf{V}^\top \mathbf{x} = \mathbf{1}\}) \subseteq C_V$, as claimed. \square

For design matrices of the form (8), the claim in the theorem follows directly from the next lemma.

LEMMA 15. *There exists a homomorphism from the face lattice of C_V to the face lattice of C_A that associates to each face of C_V of dimension d_F the (unique) face of C_A of dimension $m + d_F$ containing it.*

PROOF. Instead of concerning ourselves with C_V , we find it convenient to deal with the $d - m$ -dimensional polyhedron in \mathbb{R}^d

$$(15) \quad T_V = C_A \cap \{\mathbf{t} = (t_1, \dots, t_d)^\top \in \mathbb{R}^d : t_j = 1, j = d - m + 1, \dots, d\}.$$

In light of the next result, T_V and C_V have the same combinatorial properties.

LEMMA 16. *The polyhedra T_V and C_V are combinatorially equivalent.*

PROOF. By Lemma 14, $\mathbf{z} \in C_V$ if and only if $(\frac{\mathbf{z}}{\mathbf{1}}) \in T_V$. Thus the coordinate projection map $\pi : \mathbb{R}^d \rightarrow \mathbb{R}^{d-m}$ given by $\pi(x_1, \dots, x_d) = (x_1, \dots, x_{d-m})$ defines a bijection between C_V and T_V . Since π is a linear mapping, C_V and T_V are affinely equivalent, hence combinatorially equivalent. \square

It follows from Lemma 16 that there exists a bijection between C_V and T_V that is also a bijection between boundary points of C_V and points on the relative boundary of T_V in such a way that the face lattices of C_V and T_V are identical. Note also that isomorphic faces of the polyhedra have the same dimension. Therefore, it is sufficient to prove that the claim of the theorem holds for T_V instead of C_V .

Using the \mathcal{H} -representation [see, e.g., Ziegler (1995), Schrijver (1998)] we write

$$(16) \quad C_A = \{\mathbf{t} \in \mathbb{R}^d : \mathbf{C}\mathbf{t} \leq \mathbf{0}\}$$

for some matrix C , where we can assume that no inequality is redundant. In particular, any face F of C_A of co-dimension k can be written as

$$\{\mathbf{t} : C\mathbf{t} \leq \mathbf{0}, (\mathbf{c}_j, \mathbf{t}) = 0, j = 1, \dots, k\},$$

where $(\mathbf{c}_1, \dots, \mathbf{c}_k)$ are the k rows of C that define the k supporting hyperplanes whose intersection with C_A is precisely F . Define

$$T = \begin{bmatrix} 0 & I_m \end{bmatrix},$$

where 0 is the $m \times (d - m)$ matrix of zeros, and I_m is the $m \times m$ identity matrix. Thus, T_V is the set of points in \mathbb{R}^d given by $\{\mathbf{t} : D\mathbf{t} \leq \mathbf{b}\}$, with

$$D = \begin{bmatrix} C' \\ T \\ -T \end{bmatrix} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} \mathbf{0} \\ \mathbf{1} \\ -\mathbf{1} \end{bmatrix},$$

where C' is the sub-matrix of C obtained by removing the rows corresponding to inequalities that may have become redundant once the sampling constraint are enforced. These inequalities are the precisely the defining inequalities for the facets that do not intersect the affine space $\{\mathbf{t} : T\mathbf{t} = \mathbf{1}\}$. Notice that, by (A1), the dimension of T_V is equal to d minus the rank of

$$\begin{bmatrix} T \\ -T \end{bmatrix},$$

which is m . Next, any face F of T_V of co-dimension k can be written as

$$F = \{\mathbf{t} : D\mathbf{t} \leq \mathbf{b}, (\mathbf{d}_j, \mathbf{t}) = 0, j = 1, \dots, k\},$$

where $(\mathbf{d}_1, \dots, \mathbf{d}_k)$ are the k rows of C' that define the k supporting hyperplanes of F . Since the points in F satisfy all the inequalities (16), it follows that F is contained in the set $F' = \{\mathbf{t} : C\mathbf{t} \leq \mathbf{0}, (\mathbf{d}_j, \mathbf{t}) = 0, j = 1, \dots, k\}$, which is a face of C_A of co-dimension k . It is also immediate to see that F' is the smallest such face. Furthermore, if G is a different face of T_V of co-dimension k , it is defined by a different set of equalities, so it is contained in a different face of C_A (of co-dimension k). If G is instead of co-dimension $k' > k$ and is also a face of F , then, $G = \{\mathbf{t} : D\mathbf{t} \leq \mathbf{b}, (\mathbf{d}_j, \mathbf{t}) = 0, j = 1, \dots, k, \dots, k'\}$, so that G is contained in the set $\{\mathbf{t} : C\mathbf{t} \leq \mathbf{0}, (\mathbf{d}_j, \mathbf{t}) = 0, j = 1, \dots, k'\}$, which is a face of C_A of co-dimension k' and also a face of F' .

Therefore, the mapping that associates to each face of T_V the smallest face of C_A containing it (and of the same co-dimension) is a lattice homomorphism from the face lattice of T_V to the face lattice of C_A . Furthermore, since the homomorphism just described is between faces of the same co-dimension, and $\dim(T_V) = d - m$ while $\dim(C_A) = d$, each face of T_V of dimension d_F is mapped to a face of C_A of dimension $m + d_F$. \square

Thus far we have assumed that the design matrix A is of full rank and has the form specified by equation (8). Now let A' be any design matrix with row span \mathcal{M} , not necessarily of the form (8), or not even of full rank. Then, $C_{A'}$ is also a polyhedral cone of dimension d , though its ambient dimension may be larger. As A' and A have the same null space, the cones $C_{A'}$ and C_A are affinely isomorphic, hence combinatorially equivalent. Thus, $\mathbf{t}' = (A')^\top x \in \text{ri}(C_{A'})$ if and only if $\mathbf{t} = A^\top \mathbf{x}$, which shows that the theorem holds for any generic design matrix A . \square

PROOF OF THEOREM 9. We show that, under both Poisson and product multinomial scheme, the MLE exists, is unique and is identical in both cases if and only if $\mathbf{t} = A\mathbf{n}$ is a point in the relative interior of C_A . If \mathbf{t} belongs to the relative interior of a face F , then both log-likelihood functions realize their suprema along sequences of points $\boldsymbol{\mu}_n \subset \mathcal{M}$ for which the limit $\exp(\boldsymbol{\mu}_n) = \hat{\mathbf{m}}^e$ is unique, satisfies the moment equations $\Pi_{\mathcal{M}}\mathbf{n} = \Pi_{\mathcal{M}}\hat{\mathbf{n}}$ and $\text{supp}(\hat{\mathbf{m}}) = \mathcal{F}$.

First, we consider the problem of maximizing the log-likelihood $\ell^P(\boldsymbol{\mu}) = (\mathbf{n}, \boldsymbol{\mu}) - \sum_{i \in \mathcal{I}} \exp(\boldsymbol{\mu}(i))$ under Poisson sampling scheme. Suppose $\mathbf{t} = A^\top \mathbf{n}$ lies inside the relative interior of a proper face F of C_A with corresponding facial set \mathcal{F} . Then, there exists a $\mathbf{z}_F \in \text{kernel}(A) = \mathcal{M}^\perp$ such that the vector $\mathbf{x}_F = \mathbf{n} + \mathbf{z}_F$ satisfies $\mathbf{t} = A^\top \mathbf{x}_F$ and $\text{supp}(\mathbf{n} + \mathbf{z}_F) = \mathcal{F}$. Furthermore, since, for any $\boldsymbol{\mu} \in \mathcal{M}$, $(\mathbf{z}_F, \boldsymbol{\mu}) = 0$, $\ell^P(\boldsymbol{\mu}) = (\mathbf{x}_F, \boldsymbol{\mu}) - \sum_{i \in \mathcal{I}} \exp(\boldsymbol{\mu}(i))$.

Define $\ell_{\mathcal{F}}^P$ and $\ell_{\mathcal{F}^c}^P$ to be the restriction of ℓ^P on $\pi_{\mathcal{F}}(\mathcal{M})$ and $\pi_{\mathcal{F}^c}(\mathcal{M})$, respectively. Explicitly,

$$\ell_{\mathcal{F}}^P(\boldsymbol{\mu}) = (\mathbf{x}_F, \pi_{\mathcal{F}}(\boldsymbol{\mu})) - \sum_{i \in \mathcal{F}} \exp(\boldsymbol{\mu}(i)) = (\mathbf{x}_F, \boldsymbol{\mu}) - \sum_{i \in \mathcal{F}} \exp(\boldsymbol{\mu}(i))$$

and $\ell_{\mathcal{F}^c}^P(\boldsymbol{\mu}) = -\sum_{i \in \mathcal{F}^c} \exp(\boldsymbol{\mu}(i))$. Therefore, $\ell^P(\boldsymbol{\mu}) = \ell_{\mathcal{F}}^P(\boldsymbol{\mu}) + \ell_{\mathcal{F}^c}^P(\boldsymbol{\mu})$. On $\pi_{\mathcal{F}}(\mathcal{M})$, the function $\ell_{\mathcal{F}}^P$ is bounded from above, continuous and strictly concave, so it is maximized by the unique point $\boldsymbol{\mu}_{\mathcal{F}}^* \in \pi_{\mathcal{F}}(\mathcal{M})$ that satisfy the first order optimality conditions on the differential of $\ell_{\mathcal{F}}^P$ [see Haberman (1974), Chapter 2] given by

$$(17) \quad (\boldsymbol{\lambda}_{\mathcal{F}}, \exp(\boldsymbol{\mu}_{\mathcal{F}}^*)) = (\boldsymbol{\lambda}_{\mathcal{F}}, \pi_{\mathcal{F}}(\mathbf{x}_F)) = (\boldsymbol{\lambda}_{\mathcal{F}}, \mathbf{n}) \quad \forall \boldsymbol{\lambda}_{\mathcal{F}} \in \pi_{\mathcal{F}}(\mathcal{M}),$$

where the second equality holds since $\mathbf{x}_F \in \mathcal{M}^\perp$ and $\text{supp}(\mathbf{x}_F) = \mathcal{F}$.

On the other hand, on $\pi_{\mathcal{F}^c}(\mathcal{M})$, the function $\ell_{\mathcal{F}^c}^P$ is negative and strictly decreasing in each coordinate of its argument. Thus,

$$\sup_{\boldsymbol{\mu} \in \mathcal{M}} \ell_{\mathcal{F}}^P(\boldsymbol{\mu}) \leq \sup_{\boldsymbol{\mu}_{\mathcal{F}} \in \pi_{\mathcal{F}}(\mathcal{M})} \ell_{\mathcal{F}}^P(\boldsymbol{\mu}_{\mathcal{F}}) = \ell_{\mathcal{F}}^P(\boldsymbol{\mu}_{\mathcal{F}}^*).$$

We now show that the above inequality is in fact an equality by finding a sequence $\{\boldsymbol{\mu}_n\} \subset \mathcal{M}$ such that

$$\lim_n \ell^P(\boldsymbol{\mu}_n) = \ell_{\mathcal{F}}^P(\boldsymbol{\mu}_{\mathcal{F}}^*).$$

To this end, let $\boldsymbol{\mu}^*$ be any vector in \mathcal{M} such that $\pi_{\mathcal{F}}(\boldsymbol{\mu}^*) = \boldsymbol{\mu}_{\mathcal{F}}^*$. Next, since \mathcal{F} is a facial set, there exists a sequence $\{\boldsymbol{\gamma}_n\} \subset \mathcal{M}$ such that:

- (i) if $i \in \mathcal{F}$, then $\gamma_n(i) = 0$, for all n ;
- (ii) if $i \in \mathcal{F}^c$, then $\gamma_n(i) < 0$ for all n and $\lim_n \gamma_n(i) = -\infty$ (the rate at which these series diverge to infinity being arbitrarily fast).

Define the sequence $\{\boldsymbol{\mu}_n\} \subset \mathcal{M}$ as $\boldsymbol{\mu}_n = \boldsymbol{\mu}^* + \boldsymbol{\gamma}_n$. Then,

$$\lim_n \boldsymbol{\mu}_n(i) = \begin{cases} \boldsymbol{\mu}^*(i) & \text{if } i \in \mathcal{F}, \\ -\infty & \text{if } i \in \mathcal{F}^c, \end{cases}$$

from which it follows that

$$\begin{aligned} \lim_n \ell^P(\boldsymbol{\mu}_n) &= \lim_n \ell_{\mathcal{F}}^P(\pi_{\mathcal{F}}(\boldsymbol{\mu}_n)) + \lim_n \ell_{\mathcal{F}^c}^P(\pi_{\mathcal{F}^c}(\boldsymbol{\mu}_n)) \\ &= \ell_{\mathcal{F}}^P(\boldsymbol{\mu}_{\mathcal{F}}^*) + \lim_n \ell_{\mathcal{F}^c}^P(\pi_{\mathcal{F}^c}(\boldsymbol{\mu}_n)) = \ell_{\mathcal{F}}^P(\boldsymbol{\mu}_{\mathcal{F}}^*), \end{aligned}$$

as desired, since

$$\lim_n \ell_{\mathcal{F}^c}^P(\pi_{\mathcal{F}^c}(\boldsymbol{\mu}_n)) = \sum_{i \in \mathcal{F}^c} \lim_n \exp(\boldsymbol{\mu}_n(i)) = 0.$$

Set $\widehat{\mathbf{m}}^e = \lim_n \exp(\boldsymbol{\mu}_n)$, and notice that $\widehat{\mathbf{m}}^e$ is the unique vector in $\mathbb{R}^{\mathcal{I}}$ such that

$$\begin{cases} \pi_{\mathcal{F}}(\widehat{\mathbf{m}}^e) = \exp(\boldsymbol{\mu}_{\mathcal{F}}^*), \\ \pi_{\mathcal{F}^c}(\widehat{\mathbf{m}}^e) = \mathbf{0}, \end{cases}$$

where uniqueness stems from the uniqueness of $\boldsymbol{\mu}_{\mathcal{F}}^*$ (it is clear that, while $\widehat{\mathbf{m}}^e$ is unique, the sequence $\{\boldsymbol{\mu}_n\}$ is not). Furthermore, $\widehat{\mathbf{m}}^e$ is random, as it depends on the facial set \mathcal{F} associated to the face of C_{Λ} exposed by $\mathbf{t} = \mathbf{A}^{\top} \mathbf{n}$. Finally, in virtue of the fact that $\text{supp}(\mathbf{n}) \subseteq \mathcal{F}$, we see that, for any $\boldsymbol{\lambda} \in \mathcal{M}$,

$$(\boldsymbol{\lambda}, \widehat{\mathbf{m}}^e) = (\boldsymbol{\lambda}_{\mathcal{F}}, \exp(\boldsymbol{\mu}_{\mathcal{F}}^*)) \quad \text{and} \quad (\boldsymbol{\lambda}_{\mathcal{F}}, \mathbf{n}) = (\boldsymbol{\lambda}, \mathbf{n})$$

so that, using (17), $\widehat{\mathbf{m}}^e$ can be characterized as the unique point in \overline{M} such that

$$(\boldsymbol{\lambda}, \widehat{\mathbf{m}}^e) = (\boldsymbol{\lambda}, \mathbf{n}) \quad \forall \boldsymbol{\lambda} \in \mathcal{M},$$

or, equivalently,

$$(18) \quad \mathbf{A}^{\top} \widehat{\mathbf{m}}^e = \mathbf{A}^{\top} \mathbf{n} \quad \text{or} \quad \Pi_{\mathcal{M}} \widehat{\mathbf{m}}^e = \Pi_{\mathcal{M}} \mathbf{n}.$$

If we instead want to maximize the log-likelihood function ℓ^M under product multinomial sampling, we need to consider only the points $\widetilde{\boldsymbol{\mu}}$ inside $\widetilde{\mathcal{M}}$ as in equation (4). Fortunately, this restriction is inconsequential. First note that, by (18) and because $\mathcal{N} \subset \mathcal{M}$, the limit $\boldsymbol{\mu}^*$ satisfies the constraints $\{(\boldsymbol{\chi}_j, \exp(\boldsymbol{\mu}^*)) = N_j, j = 1, \dots, r\}$. Next, since ℓ^M and ℓ^P differ by a con-

stant on $\widetilde{\mathcal{M}}$ and $\widetilde{\mathcal{M}} \subset \mathcal{M}$, we have that

$$\ell^M(\boldsymbol{\mu}^*) = \sup_{\tilde{\boldsymbol{\mu}} \in \widetilde{\mathcal{M}}} \ell^M(\tilde{\boldsymbol{\mu}}).$$

We conclude that the log-likelihood functions under both the Poisson and product multinomial model must have the same maximizer $\widehat{\mathbf{m}}$.

Finally, we note that if $\mathbf{t} \in \text{ri}(C_A)$, so that $\mathcal{F} = \mathcal{I}$, the arguments simplify. Explicitly, there exists a point $\boldsymbol{\mu}^* \in \widetilde{\mathcal{M}} \subset \mathcal{M}$ such that

$$\begin{aligned} \sup_{\boldsymbol{\mu} \in \mathcal{M}} \ell^P(\boldsymbol{\mu}) &= \ell^P(\boldsymbol{\mu}^*), \\ \sup_{\tilde{\boldsymbol{\mu}} \in \widetilde{\mathcal{M}}} \ell^M(\tilde{\boldsymbol{\mu}}) &= \ell^M(\boldsymbol{\mu}^*) \end{aligned}$$

which we can obtain as the unique point $\widehat{\mathbf{m}}^e \in \overline{M}$ with $\text{supp}(\widehat{\mathbf{m}}^e) = \mathcal{I}$ satisfying (18). \square

Acknowledgment. The authors would like to thank the Associate Editor for invaluable help and support.

REFERENCES

- AGRESTI, A. (2002). *Categorical Data Analysis*, 2nd ed. Wiley, New York. [MR1914507](#)
- AICKIN, M. (1979). Existence of MLEs for discrete linear exponential models. *Ann. Inst. Statist. Math.* **31** 103–113. [MR0541956](#)
- BARNDORFF-NIELSEN, O. (1978). *Information and Exponential Families in Statistical Theory*. Wiley, Chichester. [MR0489333](#)
- BIRCH, M. W. (1963). Maximum likelihood in three-way contingency tables. *J. Roy. Statist. Soc. Ser. B* **25** 220–233. [MR0168065](#)
- BISHOP, Y. M. M., FIENBERG, S. E. and HOLLAND, P. W. (1975). *Discrete Multivariate Analysis*. MIT Press, Cambridge, MA. Reprinted by Springer (2007). [MR0381130](#)
- BROWN, L. D. (1986). *Fundamentals of Statistical Exponential Families with Applications in Statistical Decision Theory*. Institute of Mathematical Statistics Lecture Notes—Monograph Series **9**. IMS, Hayward, CA. [MR0882001](#)
- ČENCOV, N. N. (1982). *Statistical Decision Rules and Optimal Inference*. Translations of Mathematical Monographs **53**. Amer. Math. Soc., Providence, RI. [MR0645898](#)
- CHRISTENSEN, R. (1997). *Log-Linear Models and Logistic Regression*, 2nd ed. Springer, New York. [MR1633357](#)
- CSISZÁR, I. (1975). *I*-divergence geometry of probability distributions and minimization problems. *Ann. Probab.* **3** 146–158. [MR0365798](#)
- CSISZÁR, I. (1989). A geometric interpretation of Darroch and Ratcliff’s generalized iterative scaling. *Ann. Statist.* **17** 1409–1413. [MR1015161](#)
- CSISZÁR, I. and MATÚŠ, F. (2001). Convex cores of measures on \mathbb{R}^d . *Studia Sci. Math. Hungar.* **38** 177–190. [MR1877777](#)
- CSISZÁR, I. and MATÚŠ, F. (2003). Information projections revisited. *IEEE Trans. Inform. Theory* **49** 1474–1490. [MR1984936](#)
- CSISZÁR, I. and MATÚŠ, F. (2005). Closures of exponential families. *Ann. Probab.* **33** 582–600. [MR2123202](#)

- CSISZÁR, I. and MATÚŠ, F. (2008). Generalized maximum likelihood estimates for exponential families. *Probab. Theory Related Fields* **141** 213–246. [MR2372970](#)
- DARROCH, J. N. and RATCLIFF, D. (1972). Generalized iterative scaling for log-linear models. *Ann. Math. Statist.* **43** 1470–1480. [MR0345337](#)
- DOBRA, A. and MASSAM, H. (2010). The mode oriented stochastic search (MOSS) algorithm for log-linear models with conjugate priors. *Stat. Methodol.* **7** 240–253. [MR2643600](#)
- DOBRA, A., FIENBERG, S. E., RINALDO, A., SLAVKOVIC, A. and ZHOU, Y. (2009). Algebraic statistics and contingency table problems: Log-linear models, likelihood estimation, and disclosure limitation. In *Emerging Applications of Algebraic Geometry* (PUTINAR, M. and SULLIVAN, S., eds.). *IMA Vol. Math. Appl.* **149** 63–88. Springer, New York. [MR2500464](#)
- DRTON, M., STURMFELS, B. and SULLIVANT, S. (2009). *Lectures on Algebraic Statistics. Oberwolfach Seminars* **39**. Birkhäuser, Basel. [MR2723140](#)
- ERIKSSON, N., FIENBERG, S. E., RINALDO, A. and SULLIVANT, S. (2006). Polyhedral conditions for the nonexistence of the MLE for hierarchical log-linear models. *J. Symbolic Comput.* **41** 222–233. [MR2197157](#)
- EROSHEVA, E. A., FIENBERG, S. E. and JOUTARD, C. (2007). Describing disability through individual-level mixture models for multivariate binary data. *Ann. Appl. Stat.* **1** 502–537. [MR2415745](#)
- EWALD, G. (1996). *Combinatorial Convexity and Algebraic Geometry. Graduate Texts in Mathematics* **168**. Springer, New York. [MR1418400](#)
- FIENBERG, S. E. and RINALDO, A. (2007). Three centuries of categorical data analysis: Log-linear models and maximum likelihood estimation. *J. Statist. Plann. Inference* **137** 3430–3445. [MR2363267](#)
- FIENBERG, S. E. and RINALDO, A. (2012). Maximum likelihood estimation in log-linear models—supplementary material. Technical report, Carnegie Mellon Univ. Available at <http://www.stat.cmu.edu/~arinaldo/Fienberg.Rinaldo.Supplementary-Material.pdf>.
- FORSTER, J. (2004). Bayesian inference for Poisson and multinomial log-linear models. Technical report, School of Mathematics, Univ. Southampton.
- FULTON, W. (1993). *Introduction to Toric Varieties. Annals of Mathematics Studies* **131**. Princeton Univ. Press, Princeton, NJ. [MR1234037](#)
- GAWRILOW, E. and JOSWIG, M. (2000). Polymake: A framework for analyzing convex polytopes. In *Polytopes—Combinatorics and Computation (Oberwolfach, 1997)*. *DMV Sem.* **29** 43–73. Birkhäuser, Basel. [MR1785292](#)
- GEIGER, D., MEEK, C. and STURMFELS, B. (2006). On the toric algebra of graphical models. *Ann. Statist.* **34** 1463–1492. [MR2278364](#)
- GEYER, C. J. (2009). Likelihood inference in exponential families and directions of recession. *Electron. J. Stat.* **3** 259–289. [MR2495839](#)
- HABERMAN, S. J. (1974). *The Analysis of Frequency Data*. Univ. Chicago Press, Chicago, IL. [MR0408098](#)
- HABERMAN, S. J. (1977). Log-linear models and frequency tables with small expected cell counts. *Ann. Statist.* **5** 1148–1169. [MR0448675](#)
- KING, R. and BROOKS, S. P. (2001). Prior induction in log-linear models for general contingency table analysis. *Ann. Statist.* **29** 715–747. [MR1865338](#)
- KOEHLER, K. J. (1986). Goodness-of-fit tests for log-linear models in sparse contingency tables. *J. Amer. Statist. Assoc.* **81** 483–493. [MR0845887](#)
- LANG, J. B. (2004). Multinomial–Poisson homogeneous models for contingency tables. *Ann. Statist.* **32** 340–383. [MR2051011](#)

- LANG, J. B. (2005). Homogeneous linear predictor models for contingency tables. *J. Amer. Statist. Assoc.* **100** 121–134. [MR2156823](#)
- LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series* **17**. Oxford Univ. Press, New York. [MR1419991](#)
- LETAC, G. (1992). *Lectures on Natural Exponential Families and Their Variance Functions. Monografías de Matemática [Mathematical Monographs]* **50**. Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro. [MR1182991](#)
- MASSAM, H., LIU, J. and DOBRA, A. (2009). A conjugate prior for discrete hierarchical log-linear models. *Ann. Statist.* **37** 3431–3467. [MR2549565](#)
- MORRIS, C. (1975). Central limit theorems for multinomial sums. *Ann. Statist.* **3** 165–188. [MR0370871](#)
- MORTON, J. (2008). Relations among conditional probabilities. Technical report. [ArXiv:0808.1149v1](#).
- NARDI, Y. and RINALDO, A. (2012). The log-linear group-lasso estimator and its asymptotic properties. *Bernoulli*. To appear.
- PACHTER, L. and STURMFELS, B. (2005). *Algebraic Statistics for Computational Biology*. Cambridge Univ. Press, New York. [MR2205865](#)
- READ, T. R. C. and CRESSIE, N. A. C. (1988). *Goodness-of-Fit Statistics for Discrete Multivariate Data. Springer Series in Statistics*. Springer, New York. [MR0955054](#)
- R Development Core Team (2005). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, available at [http:// www.R-project.org](http://www.R-project.org).
- RINALDO, A., FIENBERG, S. E. and ZHOU, Y. (2009). On the geometry of discrete exponential families with application to exponential random graph models. *Electron. J. Stat.* **3** 446–484. [MR2507456](#)
- RINALDO, A., PETROVIĆ, S. and FIENBERG, S. (2011). Maximum likelihood estimation in network models. Technical report. Available at <http://arxiv.org/abs/1105.6145>.
- ROCKAFELLAR, R. T. (1970). *Convex Analysis. Princeton Mathematical Series* **28**. Princeton Univ. Press, Princeton, NJ. [MR0274683](#)
- SCHRIJVER, A. (1998). *Theory of Linear and Integer Programming*. Wiley, New York.
- VERBEEK, A. (1992). The compactification of generalized linear models. *Statist. Neerlandica* **46** 107–142. [MR1178475](#)
- ZIEGLER, M. G. (1995). *Lectures on Polytopes*. Springer, New York. [MR1311028](#)

DEPARTMENT OF STATISTICS
MACHINE LEARNING DEPARTMENT
CYLAB
HEINZ COLLEGE
CARNEGIE MELLON UNIVERSITY
5000 FORBES AVENUE
PITTSBURGH, PENNSYLVANIA 15213
USA
E-MAIL: fienberg@stat.cmu.edu

DEPARTMENT OF STATISTICS
CARNEGIE MELLON UNIVERSITY
5000 FORBES AVENUE
PITTSBURGH, PENNSYLVANIA 15213
USA
E-MAIL: arinaldo@cmu.edu